Can Reduced Numerical Precision Improve Performance without Loss of Fidelity? A Case with A Global Ocean Model

June Wu[1,2]
Mentors: Mark Petersen[2], Luke Van Roekel[2] , Bob Robey[2] , Laura Monroe[2]
[1]University of Chicago, [2]Los Alamos National Laboratory, U.S. Department of Energy

Abstract

High-resolution global ocean-climate models are often constrained by computational resources. Typically, ocean-climate models use double precision to represent their state variable arrays. However, initial conditions and forcing data in ocean-climate models are derived from observations that are not accurate to double precision. We examine the sensitivity of the Model for Prediction Across Scales ocean component (MPAS-Ocean) to the precision of its state variables. We compare the reduced-precision model with the standard double-precision model and evaluate it against observational data. We further analyze the performance gain in terms of memory access, storage, and message passing. By using the reduced-precision floating-point format, we can redistribute saved resources towards achieving higher resolution and more robust simulations in MPAS-ocean and other ocean-climate models.

MPAS-Ocean

MPAS-Ocean is a global ocean model jointly developed by the National Center for Atmospheric Research and Los Alamos National Laboratory. MPAS-Ocean is a finite-volume model based on variable-resolution spherical centroidal Voronoi tessellations (SCVT). These variable density SCVT are capable of enhancing resolution in regions of particular interest. Hence, MPAS-Ocean is particularly well suited to regional climate simulations.

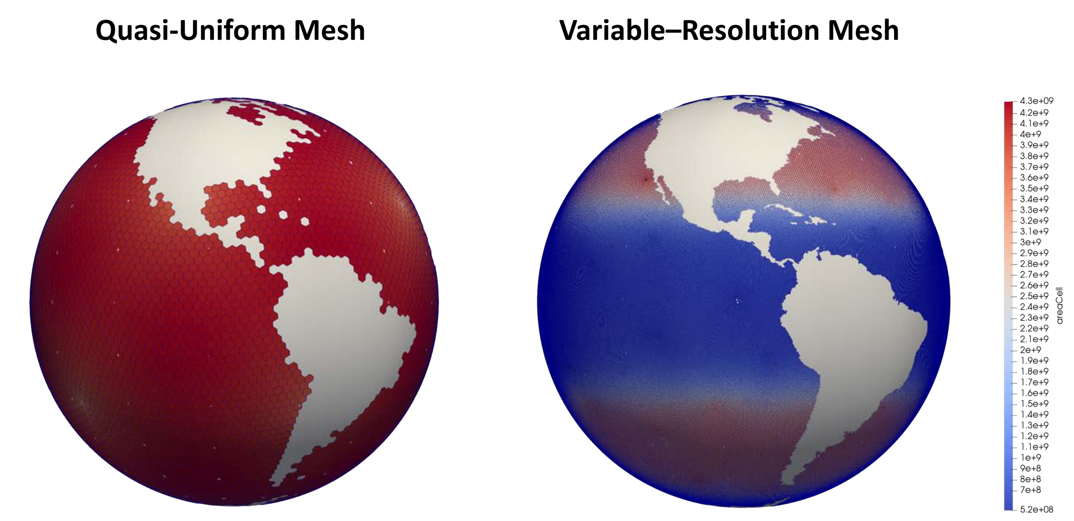

Quasi-Uniform mesh using 240 km grid cells, 7,234 horizontal cells and 100 vertical layers (QU240) .

Variable-resolution mesh from 60 to 30 km grid cells. 235160 horizontal cells and 60 vertical layers (EC60to30).

Model Versus Observations

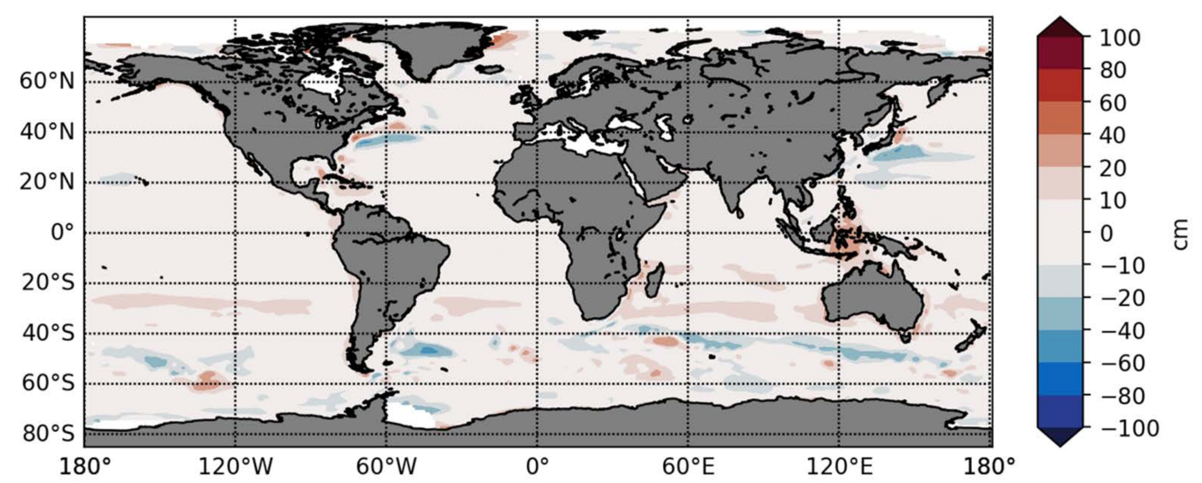

Difference in SSH averaged over one year between MPAS-Ocean and observed results.

Simulation Quality

Sea Surface Height (SSH) Averaged Over One Year

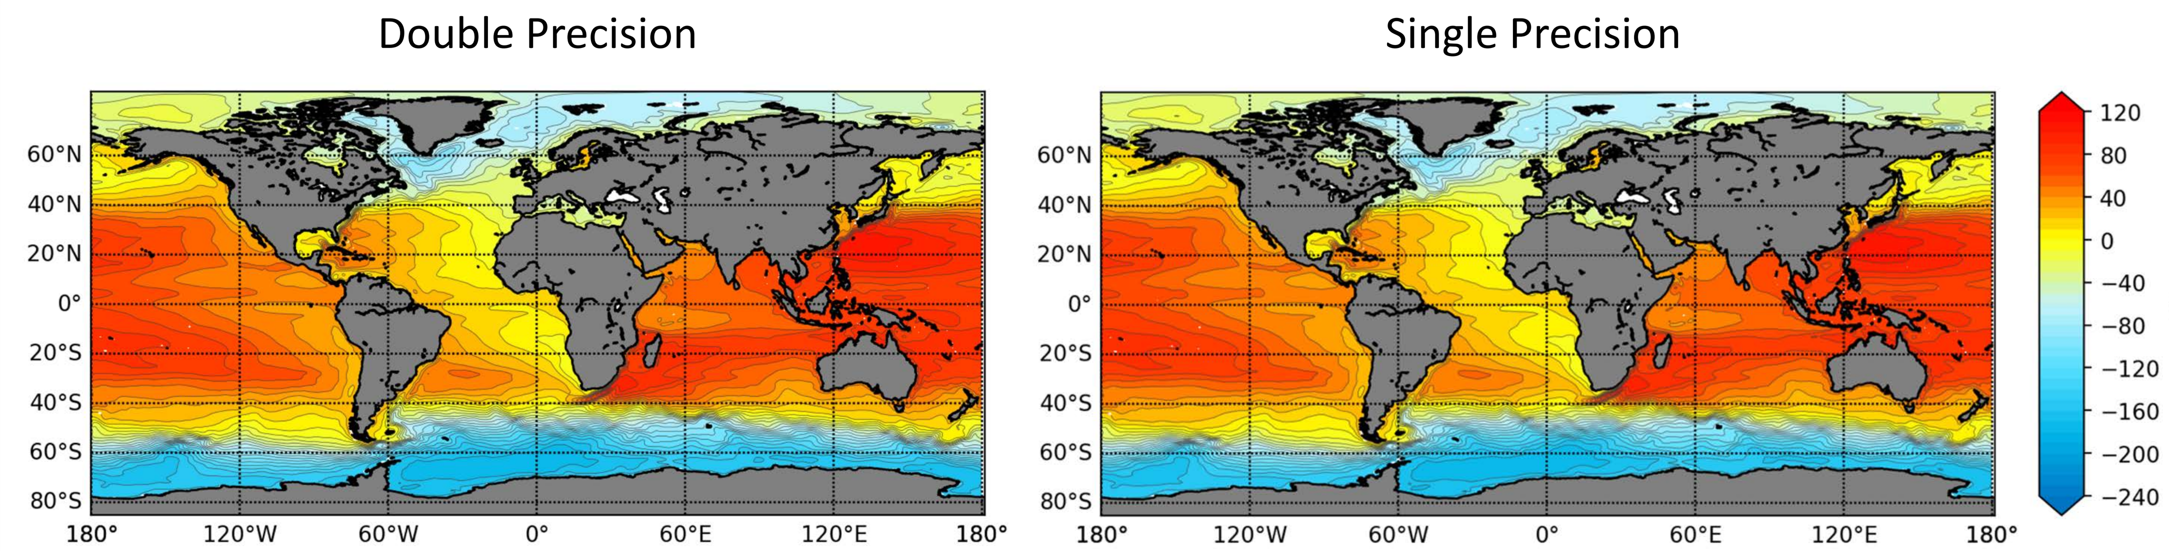

Simulated SSH averaged over one year using double-precision floating-point format in MPAS-Ocean (top left). Simulated SSH averaged over one year using single-precision floating-point format in MPAS-Ocean (top right). The difference between SSH using double precision and single precision (bottom).

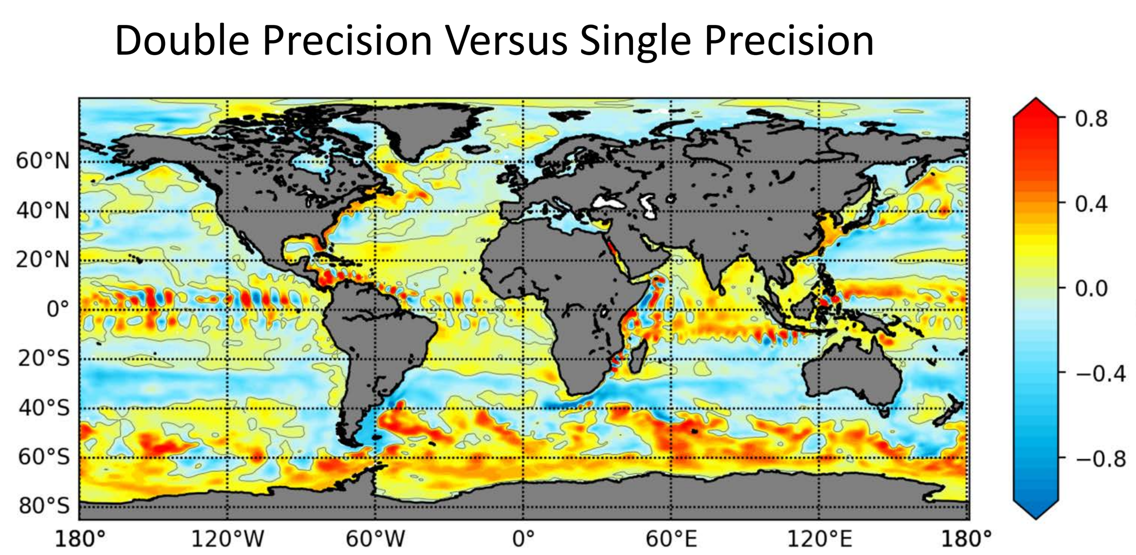

In MPAS-Ocean, four state variables are prognostic variables: layer thickness, normal velocity, salinity and temperature. In order to explore the effect of using reduced precision on MPAS-Ocean, we set prognostics variables to be single precision in one instance and double precision in another. We found no significant difference between the two levels of precision in terms of the simulation quality.

Numerical Stability

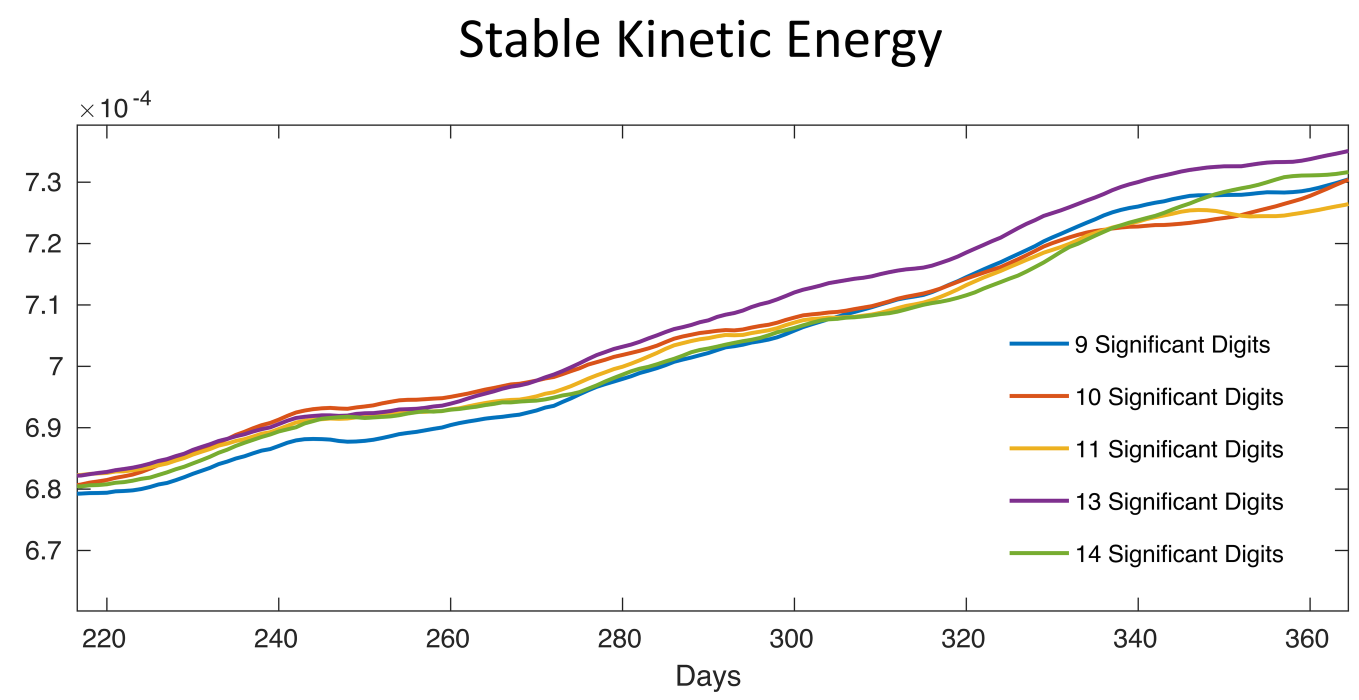

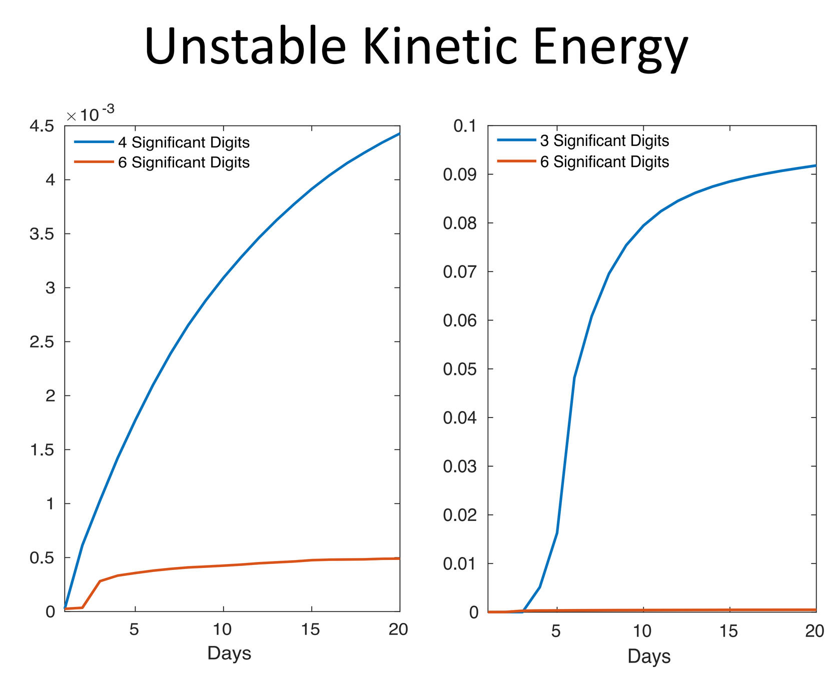

Average kinetic energy within a simulation plotted over time.

In order to find the optimal (minimum) precision of the prognostic variables in MPAS-Ocean without the loss of fidelity, we set the prognostic variables to have between 3 to 15 significant digits. We choose values for the time step and the density of the mesh so that the Courant-Friedrichs-Lewy condition is satisfied in all cases. To examine the stability of model under varying levels of precision, we approximate the average kinetic energy over one year. We find that the average kinetic energy of the model using between 6 and 15 significant digits are similar. The kinetic energy becomes blows up quickly when the number of significant digits used drops below 6 and halts the model because of the NaN (not a number) error. These NaN errors appear to be caused by roundoff error, i.e. when two number get very close, there are not enough significant digits to tell them apart when taking the difference.

Performance

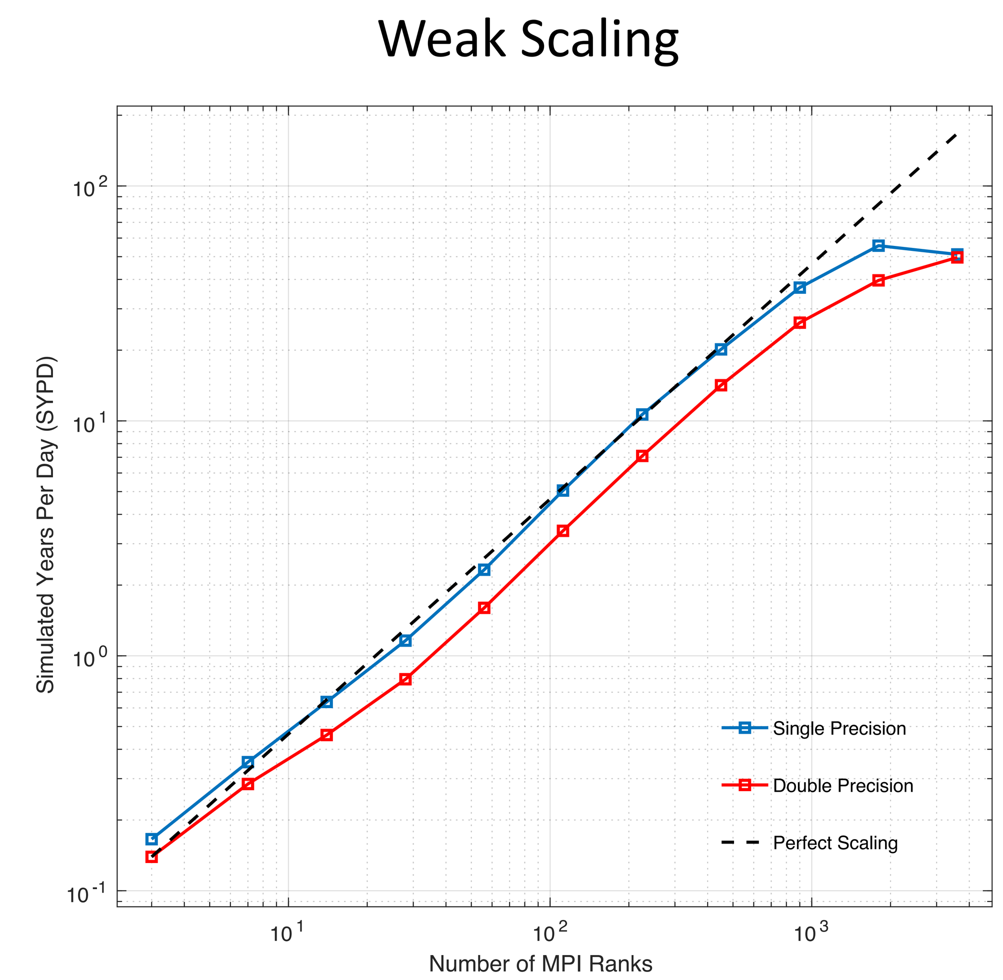

Simulated Years Per Day (SYPD) is the number of years the model can simulate within a 24-hour period, which is a measure of throughput. Without using hybrid OpenMP and MPI, the number of MPI Ranks equals to the number of cores.

All runs are tested on Grizzly, which is an institutional computing facility in Los Alamos National Laboratory. Grizzly has 53,352 cores, consisting of 1,490 nodes, each with 2 sockets, and an 18-cores Broadwell per socket. The single-precision run follows the perfect scaling more closely than the double-precision run. Before the performance tapers off, the single-precision run has approximately 1.5 times the throughput compared to that of the double-precision run.

Conclusion

We find that most of the difference between simulation quality of double precision and reduced precision are concentrated on two regions, around the equator and the Antarctic. The error around the equator is a phase shift of the equatorial Kelvin wave.

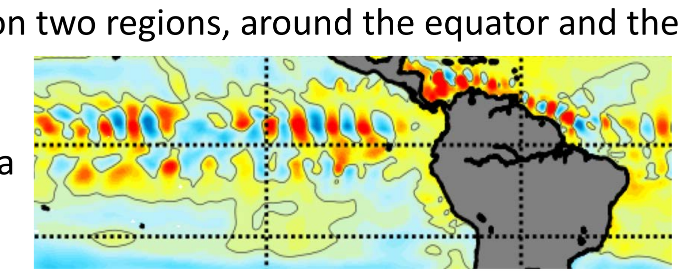

The equatorial Kelvin waves.

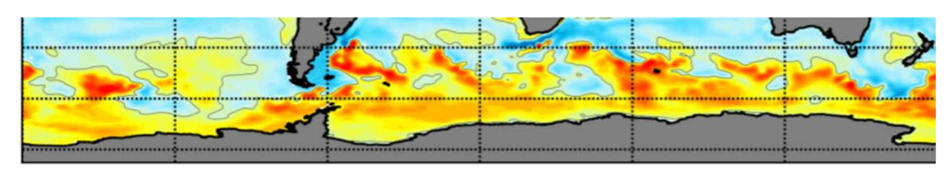

The Antarctic Circumpolar Current.

Given the reduced demand on computational resources, we could further improve simulation quality through the use of higher density mesh.

Future work may examine the effect of reduced precision within specific subroutines in MPAS-Ocean. For example, we could identify the subroutines within MPAS-Ocean that produce NaN error when using less than 6 significant digits.

References

Ringler, T., Petersen, M. R, Higdon, R. L., Jacobsen, D., Jones, P. W., & Maltrud, M. (2013). "A multi-resolution approach to global ocean modeling." *Ocean Modelling*, 69, 211–232.

Acknowledgements

This research was supported as part of the Energy Exascale Earth System Model (E3SM) project, funded by the U.S. Department of Energy, Office of Science, Office of Biological and Environmental Research. This research used resources provided by

- the Los Alamos National Laboratory Institutional Computing Program, which is supported by the U.S. Department of Energy National Nuclear Security Administration under Contract No. DE-AC52-06NA25396.
- the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

LA-UR-18-27488

MPAS Model for Prediction Across Scales

E3SM Energy Exascale Earth System Model

U.S. DEPARTMENT OF ENERGY | Office of Science